Journal of Internet Services and Applications, 2026, 17:1, doi: 10.5753/jisa.2026.7087

Using Federated Learning to Develop Congestion Control Solutions for Intelligent Transportation Systems

Carnot Braun [**University of Campinas** | ***c255785@dac.unicamp.br***]
Daniel L. Guidoni [**Federal University of Ouro Preto** | ***guidoni@ufop.edu.br***]
Allan M. de Souza [**University of Campinas** | ***allanms@unicamp.br***]

Institute of Computing, University of Campinas (UNICAMP), Av. Albert Einstein, 1251 - Cidade Universitária, Campinas - SP, 13083-852, Brazil.

Received: 02 December 2025 • **Accepted:** 13 May 2026 • **Published:** 31 July 2026

Abstract This article presents a comprehensive analysis of the application of Federated Learning (FL) to create congestion control and management solutions in the context of Intelligent Transportation Systems and Vehicular Networks. Starting from the types of communication in vehicular networks, which can be V2V, V2I, or V2X, the article presents three architectures for implementing federated learning in the context of connected vehicles: centralized, decentralized/hierarchical (using edge servers), and peer-to-peer architectures. Considering the variables available in vehicles that can be used to create congestion control solutions and how this information can be used in machine learning models, the article discusses the advantages and disadvantages of each architecture. Furthermore, the article presents data communication (V2V, V2I, and V2X) as a key element in the coordination of machine learning models and the possibilities for sharing information/models using the communication network. Finally, the article presents a case study illustrating traffic prediction with federated learning with a decentralized architecture composed of vehicles and edge servers in the network.

Keywords: Federated Learning, Congestion Control, Federated Learning Architectures, Connected Vehicles

1 Introduction

The growth of urban road traffic in recent decades is one of the main challenges faced by large urban centers. The ease of vehicle acquisition and population growth, combined with the disorderly expansion of cities, overloads the road network and, as a direct consequence, increases congestion [Tennøy and Hagen, 2021; Choudhary *et al*., 2022]. Many vehicle traffic control systems follow a traditional and centralized approach, where a server receives, processes congestion data, and performs some decision-making [Guidoni *et al*., 2020; Ebrahimi *et al*., 2024; Sun *et al*., 2020]. However, this approach has limitations: massive data collection from road sensors, roadside cameras, or onboard sensors and cameras in vehicles creates a huge amount of data to be transmitted to the servers. In real-world scenarios, this approach faces scalability, information management, and reliability issues. Furthermore, in recent years, users of traffic management services have expressed concerns about the privacy of data collected about vehicles [AlMarshoud *et al*., 2024]. A centralized approach works well in the presence of limited data and trusted actors (roadside devices and vehicles), but is restricted by the complexity of current systems.

Congestion control is one of the most critical and important functions in Intelligent Transportation Systems (ITSs). The goal of congestion control is to identify and reduce congestion on roads, improving vehicle flow and, consequently, reducing travel times [Gong *et al*., 2023; Goyal *et al*., 2024; Jurczenia and Rak, 2022]. By detecting congestion and using information about road conditions near the vehicle, the system can suggest alternative routes with shorter travel times. Advances in processing and communication capabilities in vehicles have enabled the creation of a communication network between vehicles, called vehicular ad-hoc networks (VANETs) [de Souza *et al*., 2024]. Vehicular networks allow vehicles to collaborate in traffic monitoring, where information about road conditions, average speed, or the number of vehicles on the road can be disseminated collaboratively among vehicles. Several studies in the literature have shown that distributed and collaborative solutions can achieve similar or even superior results compared to centralized solutions due to the specific mobility and communication dynamics characteristics of these networks [Pan *et al*., 2017; Gomides *et al*., 2022; de Souza *et al*., 2020].

Despite advances in ITS, congestion control solutions using a vehicular network present limitations and challenges. The large amount of data collected by vehicles, the sharing, and the frequency of information updates to neighboring vehicles imposes restrictions on the use of these networks [Balador *et al*., 2022; Paranjothi *et al*., 2020]. Distributed processing algorithms and information sharing protocols must be efficient and dynamic. Due to the specific characteristics of vehicle mobility in urban road networks, the resulting communication network between vehicles can be disconnected, meaning vehicles may not have information about other vehicles at certain times or in certain regions of the network. In this case, it is necessary to use fixed auxiliary communication stations on the roadside, called RSUs (Road-Side Units), in conjunction with central servers [Liu *et al*., 2021; Zhao *et al*., 2025]. Another challenge of these networks is the privacy of collaboration (information exchange) between vehicles, since the vehicle is disseminating information that, for example, may be related to its geographic location [Abuarqoub *et al.*, 2023]. The latency in distributed communication can be critical in certain situations, where a vehicle's near-real-time reaction to congestion is required. The existence of different layers in these networks (vehicles, RSU, and cloud servers) and the collaboration between them requires robust, specialized, and efficient communication protocols to ensure the synchronization and consistency of disseminated information. In this context, distributed congestion control solutions must operate in complex and dynamic scenarios.

In this scenario, Federated Learning (FL) emerges as a tool for proposing solutions that address the limitations and constraints of traditional congestion control approaches in vehicular networks [Zhang *et al.*, 2025, 2024]. Federated Learning [Arbaoui *et al.*, 2024; Yuan *et al.*, 2024] allows vehicles to locally train a congestion control model using only their own data (using only the information collected by the vehicle). In this scenario, no information collected by the vehicle is disseminated throughout the network, where only the model parameters are sent to an aggregator. The use of federated machine learning in vehicular networks has the potential to reduce information traffic, protect the privacy of users (vehicles), and, at the same time, enable collaborative and distributed learning among vehicles. In vehicular applications, federated machine learning has proven to be promising and efficient for various tasks [Shen *et al.*, 2024; Aalavanthar *et al.*, 2025; Li *et al.*, 2025], such as traffic prediction, incident detection, route suggestions, among others.

Considering the context of congestion control in Intelligent Transportation Systems, this article is primarily positioned as an integrative and architecture-oriented study. Unlike a survey, this article does not aim to exhaustively cover all the literature on federated learning in vehicular networks. Unlike a tutorial, the text does not aim to present a didactic implementation sequence. Our objective is to systematically examine how Federated Learning can be deployed in vehicular environments under different communication paradigms and infrastructure assumptions. The paper develops a structured conceptual analysis of centralized, hierarchical, and peer-to-peer coordination models, relating them to the operational characteristics of vehicular networks, such as dynamic topology, latency constraints, and data distribution. The experimental case study is included as an illustrative validation of these architectural discussions, demonstrating how the proposed analytical perspective can be translated into a practical deployment scenario. In this sense, the contribution of this work lies in organizing and clarifying the architectural design space of Federated Learning for congestion control applications in vehicular systems.

This work is divided as follows. Section 2 introduces the concepts of traffic congestion, the vehicular variables related to congestion control solutions, and the different forms of communication in vehicular networks. The fundamentals of federated learning in the context of connected vehicles are presented in Section 3. Section 4 presents different federated architectures for implementing congestion solutions. A case study related to vehicular flow prediction is presented in Section 5. Finally, Section 6 discusses the conclusions and future work.

2 Congestion Control and Connected Vehicles

2.1 Congestion Control

Congestion is a process resulting from the number of vehicles and the installed capacity of the road network. Congestion can be created by individual driver behavior, such as acceleration, braking, stopping at pedestrian crossings, distance from the vehicle immediately in front, overtaking decisions, inattention, and other factors. However, congestion can also be caused by factors related to road infrastructure, such as the geographic distribution of roads, road capacity (single or dual highways), traffic lights, peak hours (urban mobility dynamics), infrastructure construction work, or accidents [Treiber and Kesting, 2013]. The combination of these factors increases traffic jams by reducing the average speed experienced by vehicles, increasing travel time, increasing fuel consumption, and, consequently, increasing pollutant emissions. Congestion can also occur due to a lack of coordination between the elements that make up the transportation system, including vehicles and control infrastructure.

There are different ways to mitigate traffic congestion. Approaches based on fixed sensing infrastructure, also known as centralized ITS, use information from sensors installed along roads, images/videos from cameras, information from smart traffic lights, inductive loops, and other methods to detect congestion. The information is collected in real time and sent for processing to a central server, where a city traffic system administration, management, and control station can adjust traffic light timings, change the direction of certain roads, recommend detours, and other tasks related to traffic management [Gordon, 2015]. This approach, when used, is common in large urban centers due to the high costs of installation, maintenance, and management of the actions and actors. Furthermore, centralized ITSs require a highly stable communication network between the devices involved in collecting information, in addition to having a single point of failure (server).

Several studies have investigated congestion detection and traffic prediction in vehicular environments, often emphasizing either centralized machine learning solutions or improvements in vehicular communication protocols. However, these research topics are typically developed separately. In this work, we adopt a different perspective by explicitly relating congestion control variables, vehicular communication models (V2V, V2I, and V2X), and Federated Learning deployment strategies within a single framework. Instead of introducing a new prediction algorithm, our goal is to examine how distinct federated coordination of centralized, hierarchical, and peer-to-peer models behave under the practical constraints of vehicular networks, such as dynamic topology, intermittent connectivity, latency requirements, and privacy preservation. By discussing these architectures and linking them to congestion control scenarios, we aim to clarify the trade-offs involved in their adoption, including scalability, robustness, communication cost, and convergence characteristics. In this sense, the contribution of this article lies in organizing and contextualizing the design space of Federated Learning for congestion control in Intelligent Transportation

Systems, making explicit the architectural implications that are often only implicitly addressed in prior work.

2.2 Monitoring Variables for Congestion Control

This section presents the variables that can be collected by vehicles and used in training federated learning models for congestion control solutions. Modern vehicles are equipped with a wide variety of embedded systems and sensors that allow regular and continuous monitoring of variables associated with vehicle mobility or driver behavior. The availability of this information is essential for congestion management and control, as it allows the characterization of traffic conditions observed by vehicles in real time. Furthermore, these variables can be used to identify critical situations or they can be used in traffic forecasting models or driver decision-making. In the context of vehicular networks, these variables can be obtained through the vehicle's On-Board Unit, embedded sensors, or through direct communication between vehicles. Table 1 presents the main variables that can be used in congestion control applications considering the scenario of Intelligent Transportation Systems.

The combination of these variables creates a knowledge base about the dynamic conditions of the road network and vehicle mobility. In federated architectures, each vehicle will use its own local dataset to train machine learning models, sharing only the model parameters. This allows privacy and efficiency in communication, as well as contributing to continuous model adaptation to the current traffic conditions. In practical implementations, each vehicle will be equipped with different sensors, but federated learning considers the aggregation of trained models using different variables. This heterogeneity contributes to the construction of robust and representative models of vehicle traffic conditions and dynamics in the context of urban mobility.

2.3 Vehicular Communication

Solutions based on the principles of intelligent transportation systems and vehicular networks can be divided according to their communication methods [Cunha *et al.*, 2016]. Approaches based only on vehicle-to-vehicle communication (V2V) utilize direct messaging between vehicles, enabling them to monitor information related to speed, travel time, distance traveled, and other factors, and share this information with neighboring vehicles. This approach leverages cooperation between vehicles and enables rapid detection of small changes in traffic flow.

Solutions based on communication between vehicles and fixed communication infrastructure (RSU) assume: (i) direct communication between the vehicle and RSU; or (ii) multi-hop communication, where relay vehicles are used to create a path between the vehicle and the RSU. However, multi-hop communication depends on reliable and robust communication protocols, as the communication network is highly dynamic. This approach is named V2I (vehicle-to-infrastructure). The RSU acts as an information aggregation point, where each RSU is responsible for a specific region of the road network to be monitored. The RSU can aggregate and process received information locally and disseminate alerts or suggest alternative routes when congestion is detected. These solutions reduce the dependence on cloud servers and improve response time compared to centralized ITS, but involve the installation costs of auxiliary communication infrastructure.

A natural evolution of the V2V and V2I approaches is V2X (vehicle-to-everything) communication, where the integration of vehicles, communication infrastructure, and smart devices creates a single and cooperative network. This model, in addition to V2V and V2I communications, includes V2P (vehicle-to-pedestrian) and V2N (vehicle-to-network) communication, including various technologies such as 5G, Wi-Fi, 6LoPan, and others. The interoperability of different technologies and multiple data processing techniques, including machine learning, enables real-time or near-real-time decision-making in urban traffic management. Machine learning techniques are used for the massive processing of large volumes of data to predict and manage urban traffic and congestion. These techniques are used to extract useful information and knowledge about the dynamics of urban mobility. Predictive models can be used to anticipate congestion formation. However, it is necessary to standardize communication protocols, the format of exchanged information, and the synchronization of multiple communicating entities, in addition to having robust security mechanisms, since many sensitive data are collected by multiple agents.

More information regarding the types of communication used in applications for the connected vehicles can be found at [Soto *et al.*, 2022].

2.4 Considerations

Despite advances in recent years, several challenges remain in massive data collection for congestion control in intelligent transportation systems. Scalability remains a major obstacle, as there is a growing number of vehicles and entities with different sensing and communication capabilities. Response time for various applications remains a critical factor, especially in situations where decision-making must be instantaneous, such as rerouting in the event of an accident. A challenge that has gained attention in recent years is data privacy, where information about the location and behavior of vehicles and drivers can be used for other purposes besides the intelligent transportation solutions. Adaptability to different urban contexts is a recurring challenge, as traffic and congestion characteristics can vary depending on the topography of the road network, driver behavior, peak times, and the availability of communication infrastructure. These challenges and limitations reinforce the need for decentralized and cooperative solutions. In this context, federated learning can serve as a promising alternative for integrating communication, privacy, and intelligence in congestion control systems.

3 Fundamentals of Federated Learning in the Vehicular Context

Federated Learning is a form of distributed learning that enables collaborative training of an artificial intelligence model

Table 1. Key in-vehicle variables for congestion control and management

Variable	Acquisition Source	Purpose in Congestion Control
Average speed	Wheel speed sensor or GPS module	Identify local flow reduction and detect emerging congestion
Acceleration / Deceleration	On-board accelerometer	Detect irregular flow behavior and anticipate sudden stops
Vehicle density	V2V neighbor counting or RSU aggregation	Estimate lane occupancy and traffic intensity on a road segment
Traffic flow (veh/min)	RSUs, in-road sensors, or V2I message counters	Assess operational capacity and locate bottlenecks
Geographical position (GPS)	GNSS receiver (GPS/GLONASS/Galileo)	Map traffic conditions spatially and support network-wide situational awareness
Travel time / Average delay	Trip history (on-board) or RSU aggregation	Compare actual vs. expected times to flag critical regions
Fuel / energy consumption	Engine management unit / BMS (EVs)	Relate efficiency to congestion level; evaluate eco-routing impacts
Driver behavior indicators	Steering/brake sensors, ADAS, in-cabin cameras	Infer driving style (e.g., harsh braking, lane changes) and stress on traffic flow
Temperature / weather conditions	On-board or external environmental sensors	Adjust predictions considering weather influence on mobility
Active safety systems status	ABS/ESC/traction control telemetry	Reveal harsh braking and elevated collision risk in dense traffic
V2V / V2I message exchange	Vehicular comms stack (DSRC/C-V2X/5G)	Measure local connectivity and support cooperative control actions
Lane occupancy / obstacles	Ultrasonic sensors, cameras, radar/lidar, V2V	Detect lane saturation and advise local redistribution of flow

without the need to centralize data in a single entity. Unlike centralized learning, federated learning is implemented using an aggregator server and participating devices. These devices, which can be vehicles or RSUs in connected vehicles scenarios, train the model independently using their local data, where only the model parameters are sent to an aggregator entity. The aggregator entity receives multiple models and, using an aggregation function, aggregates the models and sends the new model to the devices. These steps are repeated until the model reaches adequate convergence [Maroua, 2024]. A difference compared to traditional machine learning is privacy preservation, as the data remains with the clients, also reducing the amount of data transmitted over the network. This approach is suitable for the vehicular environment, where data is collected by multiple vehicles and devices where it is naturally and geographically distributed.

The application of federated learning in the vehicular context can be implemented in different ways considering the vehicular communication infrastructure. In a centralized communication architecture, an aggregator server, usually located in the cloud, is responsible for receiving, aggregating, and distributing updated models to participating devices. In a decentralized/hierarchical architecture [Chellapandi *et al.*, 2024], we introduce one or more intermediate layers between the server and the devices. An intermediate layer can be implemented in the RSU, which receives models trained by vehicles in a given region. The RSU performs an aggregation of the received models and forwards the resulting model to the aggregator server, which will receive and aggregate the models from multiple RSUs. The hierarchical architecture reduces latency and communication costs, since vehicles/devices forward their models to a processing unit at the network edge, and only one aggregated/processed model is sent to the cloud. In the peer-to-peer architecture [Posner *et al.*, 2021; Du *et al.*, 2020; Gupta *et al.*, 2024], there is no role for an aggregator server or central coordinator. Instead, vehicles disseminate their models and perform aggregation among themselves, which increases resilience but also increases the complexity of model convergence. The cross-device and cross-silo variations describe the scope of collaboration between the entities participating in the learning process. While collaboration in the cross-device approach is performed by several intermittent, heterogeneous, and distributed devices (vehicles), the cross-silo variation considers a few participants with high processing power, such as transportation or urban monitoring companies.

Federated Learning offers several advantages compared to traditional ITS approaches. The first advantage is preserving the privacy of vehicles and drivers. Data collected by vehicles, such as their geographic location, average speed, travel time, daily mobility pattern, driving patterns, and vehicle condition information, are stored locally. This prevents the exposure and leakage of this information and reduces the risk of data breaches. Another advantage is the reduction in communication task, as the massive amount of raw data collected is not sent to the server/aggregator. Instead of centralized learning, federated learning enables collaborative training in highly dynamic environments. Vehicles or RSUs can perform continuous learning based on the continuous data collected by the vehicles. Continuous learning enables adaptation when traffic conditions or vehicle flow changes. Thus, federated learning is suitable for developing applications in Intelligent Transportation Systems, especially applications related to congestion control, where data is naturally decentralized in vehicles, the environment is highly dynamic, and decisions can be made in near-real time [Chellapandi *et al.*, 2024].

Despite its numerous advantages, the use of federated learning in connected vehicles presents challenges. One challenge is the heterogeneity of hardware and connectivity among the vehicles participating in the distributed training process, making it difficult to balance information across participants. The inherently highly dynamic environment is also challenging. Data synchronization across training rounds must consider vehicles that are no longer present in the network or are temporarily disconnected due to the lack of neighboring vehicles to forward their information. The vehicular environment is also susceptible to adversarial attacks, such as model poisoning. In this attack, one or more malicious participants attempt to corrupt the global model by training local models with data that does not represent the mobility verified and experienced by the vehicle during its journey. Several studies in the literature present solutions to Intrusion Detection Systems (IDS) and privacy preservation in the connected vehicle environment [Abdel Hakeem and Kim, 2025].

4 Federated Learning Architectures for Congestion Control

4.1 System Description

The typical scenario for applying federated learning to vehicular networks for congestion control requires vehicles equipped with an On-Board Unit (OBU) with storage, processing, and communication capabilities. Thus, the vehicle is capable of collecting and storing its information, training artificial intelligence models, and collaboratively participating in the federated learning update process. The hardware heterogeneity among participating vehicles is noteworthy, ranging from simple embedded systems with limited capabilities to OBUs with significant computing power in the case of autonomous vehicles. Hardware diversity impacts training time, load balancing, and communication stability/diversity [Ye *et al.*, 2023].

Connected vehicles create a communication network that can have different characteristics, including V2V, V2I, or V2X. The availability of different communication methods directly impacts the architecture available for federated learning. Figure 1 illustrates the architectures. If V2X communication is available, a centralized architecture can be used to perform federated learning. If both V2V and V2I communication are available, a decentralized and hierarchical architecture may be the best option. However, if only V2V communication is available for exchanging information between vehicles, a peer-to-peer architecture can be employed. Based on the available communication methods, federated learning can utilize different coordination strategies. Below, we present three architectures for implementing federated learning for congestion control in the context of Intelligent Transportation Systems.

From a system perspective, the proposed environment can be understood as a set of interacting layers rather than isolated entities. At the lowest layer, vehicles continuously collect mobility-related variables through onboard sensors and communication modules. These data remain locally stored and are used to perform model updates during each federated training round. Depending on the available communication mode (V2V, V2I, or V2X), the updated model parameters are then transmitted either directly to an aggregation entity or indirectly through relay vehicles or RSUs. The aggregation layer, whether centralized, hierarchical, or fully decentralized, does not merely combine parameters, but also defines the temporal dynamics of synchronization, the frequency of communication exchanges, and the propagation of updated models back to participating nodes. As a result, communication constraints such as bandwidth availability, network fragmentation, and latency directly influence convergence speed, update consistency, and overall system responsiveness. In this sense, federated learning is not an external component added to the vehicular network, but rather a process embedded within the communication infrastructure itself, whose behavior is tightly coupled with the characteristics of the underlying vehicular environment.

4.2 Centralized Architecture using V2X Communication

The traditional federated learning architecture is centralized, consisting of two entities in the context of congestion control: vehicles and a central server. The central server is usually located in the cloud, and the vehicle must have V2X communication capabilities. In this scenario, vehicles can act as relays until they find a fixed communication unit to connect to the central server, otherwise the vehicle can use cellular communication to communicate directly with the server. In this context, each vehicle senses its own data, which can include: average speed observed on a given road, number of neighbors (density) verified along its trajectory using a V2V communication protocol, travel time, and more. Initially, the server sends an artificial intelligence model to the vehicles, and each vehicle performs a round of training with its own observed data. After training, each vehicle notifies the central server, which aggregates the models. This process repeats until the expected convergence is reached.

It is important to emphasize that sending the trained model to the central server can occur using different forms of communication. Vehicles equipped with cellular communication can send their models directly. Vehicles with only V2V and V2I communication must forward their messages using some routing protocol until they find an auxiliary communication infrastructure (RSU), to forward the trained model to a central server. The approach chosen by the vehicle has a direct impact on communication time (latency), since the network topology is highly dynamic, where pre-established communication routes may not be available. Furthermore, the dependence on a central structure directly impacts the amount of data transmitted in the communication network, and this issue is exacerbated when intermediary vehicles are needed to send model parameters to the server. It's important to note

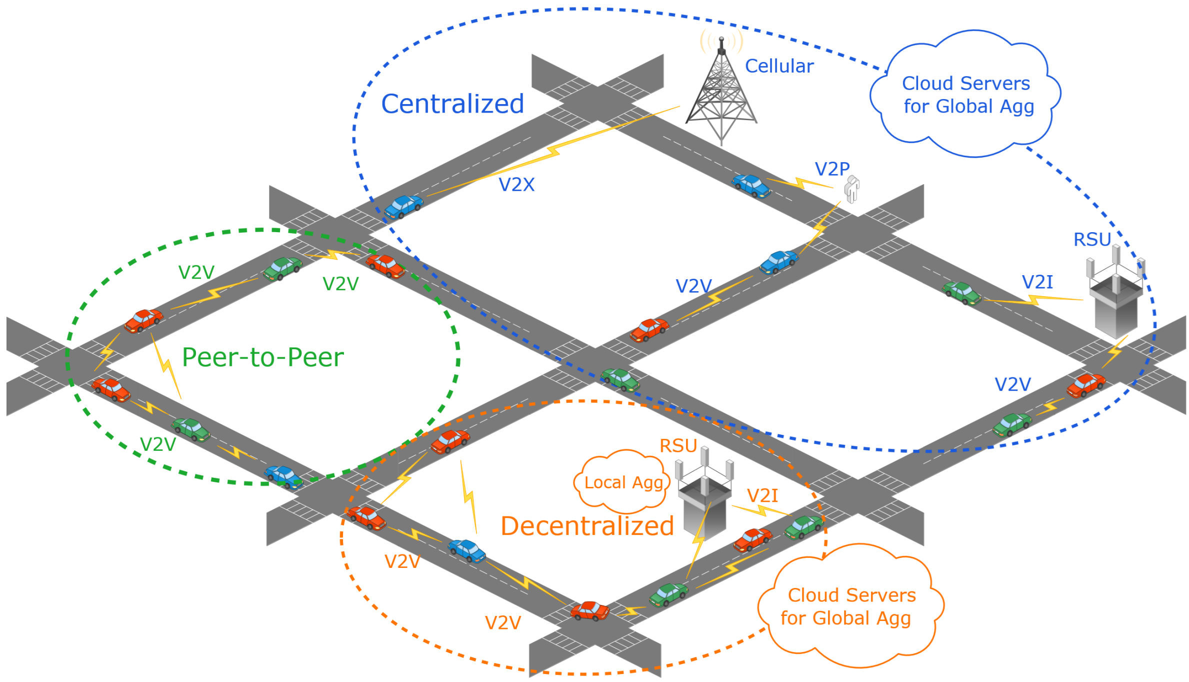

Figure 1. Communication Architectures for Federated Learning.

that in this architecture, vehicles or auxiliary communication infrastructure are used only as relays in order to create a communication route between a vehicle and the central aggregator server.

However, the greatest benefit of this architecture is its simplicity of coordination, since the central server has a global view of the network. This feature facilitates the use of techniques to improve model convergence, such as client selection, weighting techniques, and others [Lim *et al.*, 2025].

4.3 Decentralized/Hierarchical Architecture using V2V and V2I Communications

Another possible architecture for federated learning is a distributed or hierarchical architecture, creating an intermediate aggregation layer. This architecture requires the use of RSUs or edge servers deployed alongside the road network with processing power [Xiao *et al.*, 2022; Zhang *et al.*, 2023]. In this configuration, vehicles using V2I communication connect to a specific RSU nearby and, after local training, send their data to that unit. In other words, the RSU acts as a local/regional aggregator. The RSU combines multiple received updates and produces an intermediate model, which can be sent to neighboring RSUs or sent to the central aggregator server if a global model is needed. This architecture creates a hierarchical federated learning model, where aggregation can be performed at different layers/levels.

Depending on the vehicle's distance from the nearest RSU, V2V communication may be necessary, and intermediate vehicles can participate in the communication process to send trained models to the RSU. The main advantage of this architecture is the proximity of training to aggregation, since updates are processed at the edge of the network. Model aggregation by the RSU also allows the creation of specialized models for the traffic conditions of a given region, since different regions of the road network are expected to have distinct characteristics regarding vehicle flow and congestion. Another advantage is related to fault tolerance. If the global server is unavailable, the aggregation process, even if partial, can continue. A direct disadvantage of this approach is the need for fixed and auxiliary infrastructure with communication and processing capacity. Depending on the size of the region under the responsibility of an RSU, the number of updates can demand significant processing power.

The creation of an intermediate aggregation layer brings new challenges, such as the need for synchronization and information exchange between neighboring RSUs. In this case, a specific communication protocol must be designed for proper model aggregation and dissemination. Another challenge is planning the location of RSUs within the road network. Regions with high vehicle flow must be designed to ensure load balancing between regions.

An alternative way to implement federated learning using a decentralized architecture is when the RSUs train regional models. In this scenario, the vehicles only perform the task of sensing and storing their information. The sensed information is sent to the RSU, which will train the model with information from vehicles in its operating region. After training, the RSU can (i) send the trained models to the vehicles or (ii) store the model locally. In (i), the vehicles will use the model to detect congestion and calculate alternative routes. If the RSU stores the model (ii), it will perform congestion inference and notify the vehicles accordingly.

4.4 Peer-to-peer Architecture using V2V Communication

In a peer-to-peer architecture, the federated learning process occurs in a fully distributed manner, without the need for an intermediate fixed infrastructure or global server to perform model aggregation [Chen *et al.*, 2025; Yu *et al.*, 2020]. Vehicles, using their own collected data, train their models. These models are then shared with neighboring vehicles, and each vehicle, upon receiving a model, performs the aggregation process. Thus, each vehicle acts as an autonomous entity, performing model training and aggregation. By executing the model training, sharing, and aggregation process, global knowledge emerges in a distributed manner, where local models converge collaboratively. This approach is especially interesting when only V2V communication is present in the vehicular network, or when the vehicle is far from the fixed communication infrastructure and there is no communication route to it.

The peer-to-peer architecture offers greater resilience and privacy in the federated learning process, as models are shared only with neighboring vehicles or in a specific target region, reducing the risk of data exposure. Resilience and fault tolerance are inherent in this fully distributed solution. However, this architecture presents additional challenges, such as secure and efficient information sharing and local detection of inconsistent models. Furthermore, it is necessary to design solutions for controlling global model convergence and neighbor selection techniques in the aggregation process. Defining the model's dissemination area is another important factor that impacts convergence and communication costs. By increasing the dissemination area to disseminate the trained model, more vehicles will receive and perform the aggregation process. Thus, V2V communication can present challenges due to the amount of disseminated information.

4.5 Considerations

The practical implementation of these three architectures described may vary depending on the characteristics of the vehicular network. Factors such as vehicle density, response time requirements, geographic size of the road network, and availability of auxiliary infrastructure impact the solution to implement federated learning. In urban scenarios with high vehicle density, choosing an architecture with one or more aggregation layers (using RSUs) may be more efficient than other architectures. This implementation facilitates the specialization of models for specific regions of the road network, where vehicle flow varies depending on their location within the urban center. This has a direct impact on congestion detection capability in Intelligent Transportation Systems scenarios. Furthermore, the use of RSUs may offer another advantage, where elements with processing power are located at the edge of the network and can perform functions such as dynamic client selection, filtering anomalous models/updates, or load balancing among aggregating entities.

An important issue is defining the regions under each RSU's responsibility, which has a direct impact on the number of updates in the federated learning process. When the road network is large and the division of regions under each RSU's responsibility is small, a hierarchical approach (multiple levels of aggregation layers) can offer advantages to create knowledge of large regions. In this case, a periodic update protocol is necessary to create distributed knowledge among RSUs. Furthermore, RSUs can exchange information with neighboring RSUs so that the boundary areas between regions do not have limited knowledge of traffic conditions, thus ensuring greater availability of information for drivers to make decisions to mitigate congestion. It is important to note that, even when creating specialized models, there is a central aggregation server that aims to create comprehensive knowledge of the macro traffic conditions on the road network, which is important for decisions that can globally impact vehicle flow and congestion control.

In a scenario where auxiliary communication and processing infrastructure is not available, the peer-to-peer federated learning approach is a viable solution. However, this approach poses even greater challenges in the context of connected vehicles, as all communication is distributed and the vehicles are responsible for the entire federated learning process. Furthermore, vehicles must be equipped with more powerful hardware to carry out the entire learning process: collecting and storing local information, training the model, disseminating the trained model to neighboring vehicles, and aggregating multiple received models (shared by neighbors). All of these steps are performed during the vehicle's journey, as traffic conditions can vary depending on the region the vehicle is located. Another important issue is the lifetime of the collected information and the received models. Depending on when the vehicle collected/received this information, which is based on a past trajectory, the information may not be useful in order to generate knowledge regarding the vehicle's current traffic flow estimation. This issue imposes challenges of estimating new congestion and alternative routes.

4.6 Research Problems

Considering the discussion presented related to centralized, decentralized/hierarchical, and peer-to-peer architectures, communication models, and participating entities in the distributed federated learning process, we can summarize the following relevant research topics related to federated learning and congestion control in intelligent transportation systems:

- Coordination between regional aggregators: proposing robust and reliable communication protocols for sharing information between neighboring RSUs to create knowledge between regions about traffic conditions.
- Multi-layered aggregators: creating models considering one or more aggregation layers. It is important to propose federated learning aggregation functions aware of the multiple aggregation levels.
- Multi-layered communication: In the case of multiple layers, defining the communication structure among them is important considering quality of service required by the congestion control application. Propose resilient communication protocols in scenarios where an RSU becomes inaccessible.
- Integration of different communication networks: creating scalable solutions in the presence of various com-

munication technologies.

- Decentralized consensus and asynchronous convergence: Develop efficient algorithms for scenarios without an aggregator server/RSU, providing model convergence considering only the communication among vehicles.
- Client selection: Propose solutions for client selection in multi-layered or peer-to-peer architectures considering the congestion control requirements.
- Opportunistic networks: Investigate the integration of decentralized federated learning with opportunistic networks, where vehicles can create temporary groups for collaboration based on a specific interest in traffic conditions.

5 Use Case: Traffic Prediction based on Level of Service

In this use case, we apply federated learning to predict traffic flow in vehicular networks. In this context, the Level of Service (LoS) functions as a standardized qualitative metric within the discipline of traffic engineering, employed to delineate the operational conditions experienced on a road/roadway. LoS is evaluated on a continuum ranging from "A" (indicative of free flow) to "F" (indicative of forced flow or system collapse), as delineated in authoritative capacity manuals such as the *Highway Capacity Manual* (HCM). The capability to prognosticate the imminent deterioration of LoS, such as a transition from state "C" to "E" within a forthcoming 15-minute interval, is thus of paramount importance for the purposes of adaptive traffic control and management.

However, traditional predictive models based on machine learning impose a significant architectural challenge: they require the centralization of mobility data. This centralized approach faces two critical barriers. First, privacy: raw flow and velocity data, when aggregated, can expose sensitive mobility patterns of drivers. Second, scalability and latency: the massive volume of telemetry data continuously generated by sensors and vehicles (V2X) overwhelms communication links and central servers. Therefore, in this scenario, FL emerges as an enabling architecture to overcome these barriers, by allowing models to be trained in vehicles or RSUs. Additionally, this approach drastically reduces the communication load, making the creation of robust, large-scale predictive models a practical and viable solution for congestion control.

In this case study, we aim to address the following research question: can a decentralized federated learning architecture achieve traffic prediction performance comparable to a centralized baseline while operating under geographically distributed data partitions typical of vehicular networks? The objective is not to maximize predictive performance, but to examine whether the architectural benefits of federated coordination can be obtained without a significant loss in predictive quality.

5.1 FL for LoS Forecasting

In this approach, we implemented a fully decentralized FL framework, in which the RSUs, distributed geographically throughout the scenario, act as independent peers. The objective is to cooperatively train a global model capable of predicting future sequences of the LoS. Each RSU is responsible for collecting time-series data from roadways within its coverage zone, using local features such as average vehicle counts and observed Level of Service. The chosen prediction model architecture is a *Recurrent Neural Network* (RNN) of the *Long Short-Term Memory* (LSTM) type, recognized for its effectiveness in capturing temporal dependencies in sequential data. The model is structured for a *sequence-to-sequence* task, receiving a sequence of 12 elapsed *steps* (seconds) as input and being trained to predict the LoS sequence for the subsequent 5, 10, or 15 *steps* (horizons). To this end, the model's output layer is composed of H units (where H is the horizon) with `softmax` activation, each predicting the probability distribution over the 6 LoS classes (A-F) for a future *step*.

The federated training methodology, implemented through a decentralized gossip-based protocol analogous to Federated Averaging, operates in a cyclical manner. Initially, each RSU conducts local model training, adjusting model weights based solely on its proprietary dataset for a specified number of epochs, such as five. This localized training phase is pivotal for ensuring privacy preservation, as raw traffic data remains within the confines of the RSU's infrastructure. Upon the conclusion of local training, each RSU communicates its revised model parameters directly to adjacent RSUs (peers). Subsequently, each RSU awaits the reception of models from these peers, thereafter executing the aggregation phase locally by computing a weighted average of the received parameters to generate an updated model. This iterative cycle is repeated over a predetermined number of rounds, facilitating the progressive convergence of models across all RSUs toward a comprehensive global solution that effectively integrates knowledge derived from diverse local traffic patterns.

5.2 Performance Analysis

5.2.1 Simulation Setup

The simulation ecosystem was built using Python, with the *Keras* (TensorFlow) library for defining and training LSTM models. The RSU communication and decentralized aggregation logic were implemented to simulate the gossip-based protocol.

The effectiveness and efficiency of the proposed solution were validated using a real-world dataset. We used the PEMS dataset California Department of Transportation (Caltrans) [2025], widely used in traffic forecasting tasks. The original data were pre-processed and geographically segmented to represent the distributed vehicular scenario. For this, we used metadata with the coordinates of each collection station, allowing us to apply the K-Means algorithm to group the points into six distinct zones, each corresponding to a RSU. In our study, the geographical delimitation of each RSU was determined by grouping the coordinates of the collection points using the K-Means algorithm. Thus, each RSU corresponds to a zone defined by the geographical proximity of the data.

To better clarify the evaluation scenario, we emphasize that

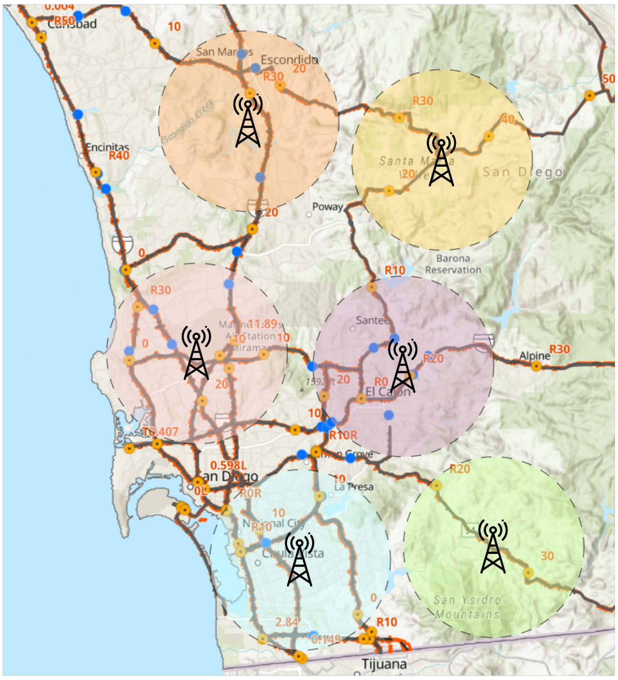

Figure 2. Distribution of the RSUs in the scenario.

the experimental setup was designed to emulate a geographically distributed vehicular environment in which multiple RSUs operate over distinct urban regions. Each RSU was associated with a subset of traffic monitoring stations obtained from the PeMS dataset, grouped according to spatial proximity. This configuration simulates realistic deployment conditions in which traffic patterns vary across different areas of a metropolitan region, as we can see in Figure 2. During the experiments, each RSU maintained its own local dataset and independently executed model training before participating in the decentralized aggregation process. The communication rounds were organized to reflect periodic model exchanges between neighboring units, without assuming a centralized coordinator. This setup allows us to evaluate not only predictive performance but also the behavior of the learning process under distributed data partitions and heterogeneous traffic conditions, which are typical characteristics of real vehicular networks. In this study, the term communication round refers to a full cycle of local training, parameter exchange, and aggregation across RSUs, while local epochs denote the number of passes over each RSU's local dataset during a given round.

After preprocessing and spatial clustering, the dataset was partitioned among the RSUs according to geographic proximity, resulting in approximately 21–24k samples per RSU (21,538; 23,550; 23,710; 23,730; 23,551), with a total of 116,079 time-series samples considered in the study. Each local dataset was divided into training (70%), validation (15%), and testing (15%) subsets following a temporal split to preserve the chronological structure of the traffic data. The LSTM model was trained using categorical cross-entropy as the loss function with the Adam optimizer and a learning rate of 0.001. Local training was performed for five epochs per communication round.

It is important to emphasize that the purpose of this case study is illustrative, serving to demonstrate the practical feasibility of the architectural discussion presented in this paper rather than to provide an exhaustive experimental benchmark. The reported performance curves correspond to the average results of 33 independent simulation runs conducted under identical configuration settings but with different random seed initializations.

Figure 3 shows the resulting LoS distribution for each of the six RSUs. The decentralized simulation was run over 200 aggregation rounds, with the participation of all clients (5 local epochs per round). It is possible to observe that traffic conditions are different in each region where RSU operates. Therefore, decentralized learning can capture variations in vehicle flow across different regions.

5.2.2 Evaluation

The system's performance evaluation focused on comparing the predictive effectiveness of the proposed Federated Learning approach against a traditional Centralized baseline across three different prediction horizons (H=5, H=10, and H=15).

Figure 4 presents the performance of the centralized model. As expected, the model achieves stable convergence. However, a clear trend is observed where predictive performance slightly degrades as the time horizon increases, with H=5 consistently outperforming H=15. This is consistent with the inherent difficulty of long-term traffic forecasting.

Figure 5 depicts the learning dynamics of the federated model over 200 communication rounds. The results exhibit a stable and robust convergence behavior. In particular, the federated model attains accuracy and recall values that are comparable to those of the centralized baseline, with performance stabilizing between approximately 50 and 75 rounds, depending on the prediction horizon. The accuracy trajectory displays a steadily increasing trend, indicating that the decentralized aggregation strategy enhances learning efficacy, especially under distributed and heterogeneous data conditions.

The federated approach maintains competitive performance even for the more challenging H=15 horizon. The recall metric, which is vital for detecting less frequent traffic states (like congestion), remains strong, suggesting that the model effectively learns from the heterogeneous data distributions of the different RSUs. Although the results indicate that the federated approach achieves performance close to the centralized baseline, this comparison should be interpreted with caution. Similar performance was obtained in a scenario where all participants remain available, they use the same model configuration, and participate in all aggregation rounds. In a real-world deployment, vehicle mobility, temporary communication failures, delays in parameter transmission, and unbalanced data distribution can reduce the stability of the training process. Furthermore, sending model parameters, while reducing the exposure of raw data, does not completely eliminate communication costs or the privacy risks associated with inference or reconstruction attacks. Thus, the results should be understood as initial evidence of architectural feasibility, and not as proof of the superiority of federated learning over centralized solutions.

The degradation observed with increasing prediction horizon also highlights a relevant limitation for congestion control applications. Longer horizons are more useful for preventive

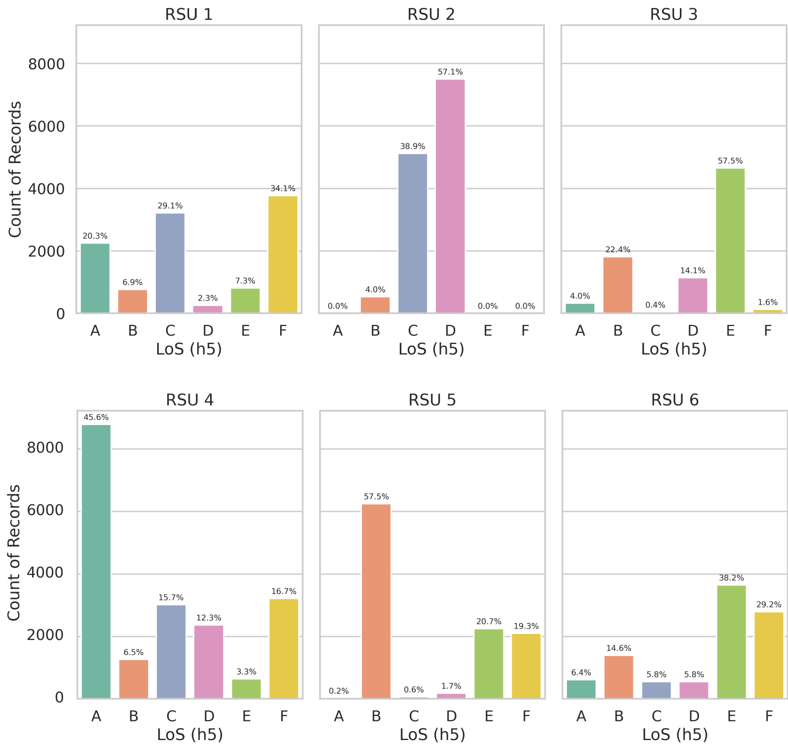

Figure 3. Distribution of LoS classes across the RSUs, highlighting regional heterogeneity in traffic conditions. The variability among zones motivates the use of decentralized learning to capture localized mobility patterns.

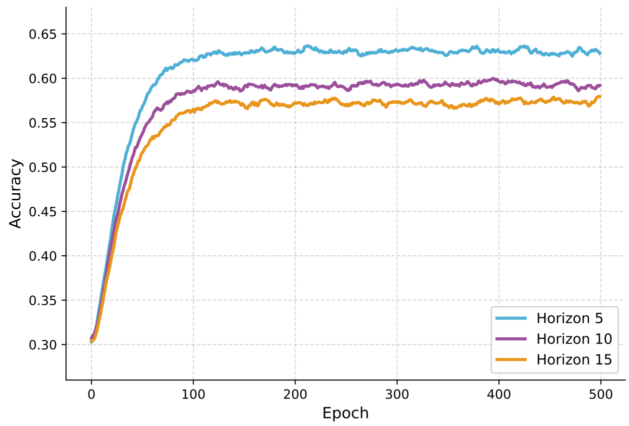

(a) Evolution of accuracy.

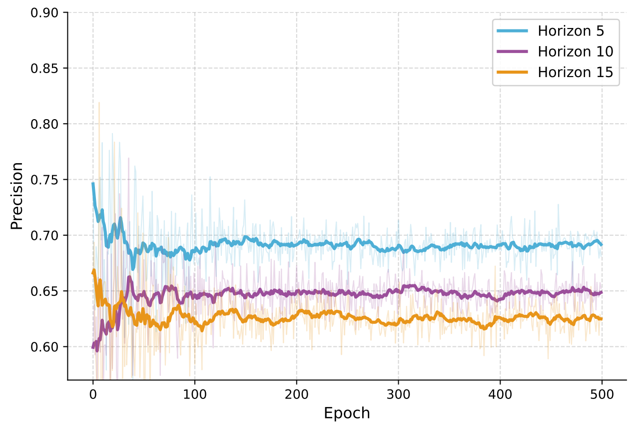

(b) Evolution of precision.

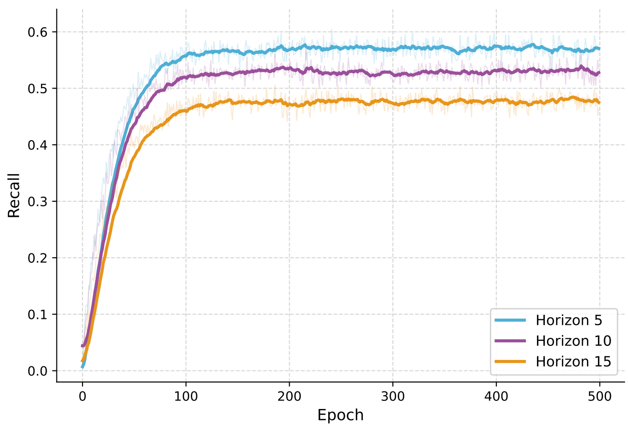

(c) Evolution of recall.

Figure 4. Performance metrics (Accuracy, Precision, Recall) for the Centralized model across prediction horizons. As expected, shorter horizons (H=5) exhibit higher predictive stability, while longer horizons introduce gradual performance degradation due to increased forecasting uncertainty.

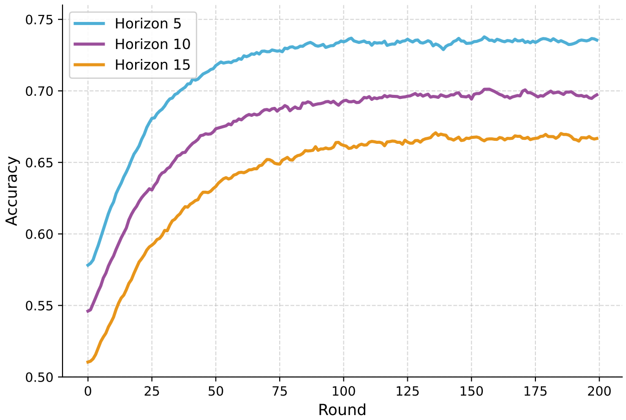

(a) Evolution of accuracy considering the rounds of federated learning.

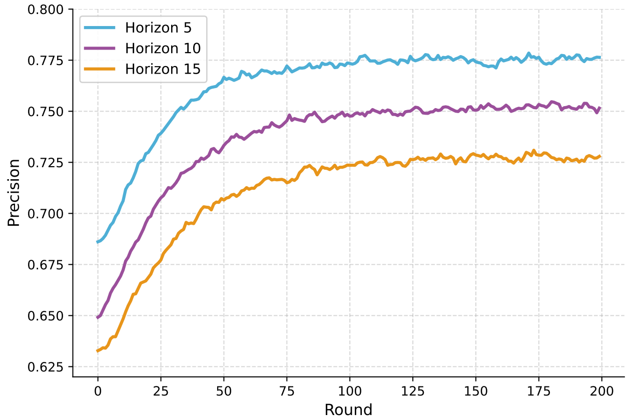

(b) Evolution of precision considering the rounds of federated learning.

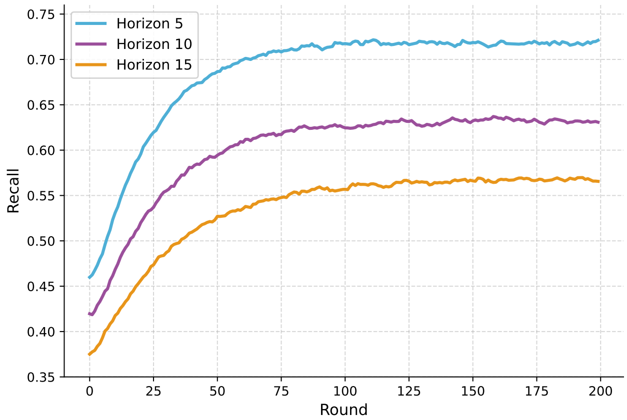

(c) Evolution of recall considering the rounds of federated learning.

Figure 5. Performance metrics (Accuracy, Precision, Recall) for the Federated model across rounds and horizons. The curves show stable improvement and eventual convergence across horizons, demonstrating that decentralized aggregation achieves performance comparable to the centralized baseline.

actions, such as early re-routing or traffic light adjustments, but they also accumulate greater temporal uncertainty. In this sense, the choice of horizon represents a compromise between operational utility and predictive reliability. For applications that demand near real-time decisions, short horizons may be more suitable; for tactical planning, longer horizons would require additional calibration mechanisms, frequent model updates, or explicit uncertainty handling.

5.3 Discussion of Assumptions and Limitations

From a deployment perspective, the fact that the federated approach achieves performance comparable to the centralized baseline has practical implications for Intelligent Transportation Systems. In real-world scenarios, transmitting raw mobility data from vehicles or RSUs to a central server can generate significant communication overhead and raise privacy concerns, especially when location traces and traffic patterns are involved. By exchanging only model parameters, federated learning can reduce the volume of transmitted data while preserving sensitive information at the source. This characteristic makes the approach attractive for large-scale ITS deployments where bandwidth, latency, and data governance constraints must be considered.

At the same time, it is important to clarify the limitations of the case study and its implications for real-world deployments. The experiments assume honest participants during aggregation, which excludes scenarios of model poisoning, sending inconsistent updates, or strategic behavior of compromised nodes. They also assume static placement of RSUs, while real urban deployments may require repositioning, gradual infrastructure expansion, or operation with partial coverage. The homogeneous configuration of the models simplifies aggregation but does not fully represent environments considering vehicles and RSUs with different computational capabilities, different sampling rates, and different connectivity conditions. Furthermore, network interruptions, asymmetric delays, and partial node participation were not explicitly modeled. These limitations do not invalidate the objective of the case study but delimit its scope: the results demonstrate the initial feasibility of the architecture, while robustness under adverse operational conditions remains a direction for future research. Future work should investigate more dynamic deployment conditions and robustness mechanisms to better approximate real operational environments.

6 Conclusion and Future Work

This article presented a conceptual and architecture-oriented analysis of the use of Federated Learning in congestion control and management solutions for Intelligent Transportation

Systems. The central contribution of this work does not lie in the proposal of a new learning algorithm, nor in an exhaustive review of the literature, but in the organization of the design space that connects vehicular communication models, aggregation entities, data distribution, and the operational requirements of urban applications. From this perspective, the article discussed three architectural alternatives: centralized, hierarchical/decentralized, and peer-to-peer, highlighting their trade-offs in terms of latency, communication cost, privacy, robustness, and convergence.

We presented the limitations of traditional approaches related to scalability, system response time, and privacy. Considering the ability of vehicles and auxiliary infrastructure at the network edge to perform distributed training of machine learning models, federated learning allows the creation of cooperative models without the need to share raw and sensitive data. Aspects related to different forms of communication in the context of connected vehicles were the starting point for the discussion of different architectures in the implementation of the federated machine learning process. In the centralized architecture, vehicles perform the task of training models and sending parameters to a central aggregator server. In the decentralized/hierarchical approach, intermediate layers can be implemented to perform training or partial aggregation of models. Finally, we present the peer-to-peer architecture, a distributed approach where vehicles exchange models with each other and perform the aggregation task. Each architecture was discussed considering the appropriate forms of communication for its implementation, and challenges from the perspective of latency, communication cost, robustness, and convergence.

This article also presents a case study describing the implementation of federated learning in vehicular networks for predicting vehicle traffic, indicating that the architectures presented allow for advancements in congestion control and vehicular data, which are naturally distributed. The use of local variables enables the creation of dynamic models that evolve and can represent the traffic conditions of the road network in urban centres. The integration of Federated Learning, edge computing, and V2X communication creates a cooperative environment for congestion control applications. Future work should focus on robust hierarchical aggregation, coordination between network service units (RSUs), and decentralized consensus, consolidating federated learning for the next generations of connected and intelligent transportation systems.

Funding

This research was partially funded by FAPEMIG (Minas Gerais Research Foundation), grant #APQ-01649-23.

Authors' Contributions

CB contributed to the conception and design of the study, developed the federated learning architectures, conducted the experiments, analyzed the results, and prepared the initial draft of the manuscript. DLG and AMS contributed to the conception and design of the study, participated in the development of the federated learning architectures, analyzed the results, and critically revised the manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Availability of data and materials

The datasets generated and/or analyzed during the current study are available from the corresponding author upon reasonable request.

References

Aalavanthar, A., S., F., Sundaramurthy, S., Cirillo, S., Solimando, G., and Polese, G. (2025). Multi-objective federated learning traffic prediction in vehicular network for intelligent transportation system. *PeerJ Computer Science*, 11:e2922. DOI: 10.7717/peerj-cs.2922.

Abdel Hakeem, S. A. and Kim, H. (2025). Advancing intrusion detection in v2x networks: A comprehensive survey on machine learning, federated learning, and edge ai for v2x security. *IEEE Transactions on Intelligent Transportation Systems*, 26(8):11137–11205. DOI: 10.1109/TITS.2025.3558849.

Abuarqoub, A., Alzu'bi, A., Hammoudeh, M., Ahmad, A., and Al-Shargabi, B. (2023). A survey on vehicular ad hoc networks security attacks and countermeasures. In *Proceedings of the 6th International Conference on Future Networks & Distributed Systems*, ICFNDS '22, page 701–707, New York, NY, USA. Association for Computing Machinery. DOI: 10.1145/3584202.3584309.

AlMarshoud, M., Sabir Kiraz, M., and H. Al-Bayatti, A. (2024). Security, privacy, and decentralized trust management in vanets: A review of current research and future directions. *ACM Comput. Surv.*, 56(10). DOI: 10.1145/3656166.

Arbaoui, M., Brahmia, M.-e.-A., Rahmoun, A., and Zghal, M. (2024). Federated learning survey: A multi-level taxonomy of aggregation techniques, experimental insights, and future frontiers. *ACM Trans. Intell. Syst. Technol.*, 15(6). DOI: 10.1145/3678182.

Balador, A., Cinque, E., Pratesi, M., Valentini, F., Bai, C., Gómez, A. A., and Mohammadi, M. (2022). Survey on decentralized congestion control methods for vehicular communication. *Vehicular Communications*, 33:100394. DOI: 10.1016/j.vehcom.2021.100394.

California Department of Transportation (Caltrans) (2025). PeMS data source. Available at:`https://dot.ca.gov/programs/traffic-operations/mpr/pems-source`. Acesso em: 31 mar. 2025.

Chellapandi, V. P., Yuan, L., Brinton, C. G., Żak, S. H., and Wang, Z. (2024). Federated learning for connected and automated vehicles: A survey of existing approaches and challenges. *IEEE Transactions on Intelligent Vehicles*, 9(1):119–137. DOI: 10.1109/TIV.2023.3332675.

Chen, D., Deng, T., Jia, J., Feng, S., and Yuan, D. (2025). Mobility-aware decentralized federated learning with joint optimization of local iteration and leader selection for vehicular networks. *Computer Networks*, 263:111232. DOI: 10.1016/j.comnet.2025.111232.

Choudhary, A., Gokhale, S. B., Kumar, P., Pradhan, C., and Sahu, S. K. (2022). Urban traffic congestion: its causes-consequences-mitigation. *Research Journal of Chemistry and Environment*. DOI: 10.25303/2612rjce1640176.

Cunha, F., Villas, L., Boukerche, A., Maia, G., Viana, A., Mini, R. A., and Loureiro, A. A. (2016). Data communication in vanets: Protocols, applications and challenges. *Ad Hoc Networks*, 44:90–103. DOI: 10.1016/j.adhoc.2016.02.017.

de Souza, A. M., Braun, T., Botega, L. C., Villas, L. A., and Loureiro, A. A. F. (2020). Safe and sound: Driver safety-aware vehicle re-routing based on spatiotemporal information. *IEEE Transactions on Intelligent Transportation Systems*, 21(9):3973–3989. DOI: 10.1109/TITS.2019.2958624.

de Souza, J. A. R., Cavalcanti, E. R., Bezerra, T. d. S., and Gomes, R. C. d. M. (2024). Vanets' research overview updated: past, present and future. *Journal of the Brazilian Computer Society*, 30(1):380–393. DOI: 10.5753/jbcs.2024.3354.

Du, Z., Wu, C., Yoshinaga, T., Yau, K.-L. A., Ji, Y., and Li, J. (2020). Federated learning for vehicular internet of things: Recent advances and open issues. *IEEE Open Journal of the Computer Society*, 1:45–61. DOI: 10.1109/OJCS.2020.2992630.

Ebrahimi, D., Sudarshan, G., and Alzhouri, F. (2024). Bnart: A novel centralized traffic management approach for autonomous vehicles. In *2024 20th International Conference on the Design of Reliable Communication Networks (DRCN)*, pages 1–8. DOI: 10.1109/DRCN60692.2024.10539171.

Gomides, T. S., De Grande, R. E., Meneguette, R. I., de Souza, F. S., and Guidoni, D. L. (2022). Predictive congestion control based on collaborative information sharing for vehicular ad hoc networks. *Computer Networks*, 211:108955. DOI: 10.1016/j.comnet.2022.108955.

Gong, T., Zhu, L., Yu, F. R., and Tang, T. (2023). Edge intelligence in intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 24(9):8919–8944. DOI: 10.1109/TITS.2023.3275741.

Gordon, R. (2015). *Intelligent Transportation Systems: Functional Design for Effective Traffic Management*. Springer Cham. DOI: 10.1007/978-3-319-14768-0.

Goyal, R., Elawadhi, O., Sharma, A., Bhutani, M., and Jain, A. (2024). Cloud-connected central unit for traffic control: interfacing sensing units and centralized control for efficient traffic management. *International Journal of Information Technology*, 16(2):841–851. DOI: 10.1007/s41870-023-01527-w.

Guidoni, D. L., Maia, G., Souza, F. S. H., Villas, L. A., and Loureiro, A. A. F. (2020). Vehicular traffic management based on traffic engineering for vehicular ad hoc networks. *IEEE Access*, 8:45167–45183. DOI: 10.1109/ACCESS.2020.2978700.

Gupta, V., Luqman, A., Chattopadhyay, N., Chattopadhyay, A., and Niyato, D. (2024). Travellingfl: Communication efficient peer-to-peer federated learning. *IEEE Transactions on Vehicular Technology*, 73(4):5005–5019. DOI: 10.1109/TVT.2023.3332898.

Jurczenia, K. and Rak, J. (2022). A survey of vehicular network systems for road traffic management. *IEEE Access*, 10:42365–42385. DOI: 10.1109/ACCESS.2022.3168354.

Li, M., Pan, X., Liu, C., and Li, Z. (2025). Federated deep reinforcement learning-based urban traffic signal optimal control. *Scientific Reports*, 15(1):11724. DOI: 10.1038/s41598-025-91966-1.

Lim, J., Kim, S., and Moon, S.-M. (2025). Convergence analysis of federated learning methods using backward error analysis. AAAI'25/IAAI'25/EAAI'25. AAAI Press. DOI: 10.1609/aaai.v39i18.34060.

Liu, L., Chen, C., Pei, Q., Maharjan, S., and Zhang, Y. (2021). Vehicular edge computing and networking: A survey. *Mob. Netw. Appl.*, 26(3):1145–1168. DOI: 10.1007/s11036-020-01624-1.

Maroua, D. (2024). A state-of-the-art on federated learning for vehicular communications. *Vehicular Communications*, 45:100709. DOI: 10.1016/j.vehcom.2023.100709.

Pan, J. S., Popa, I. S., and Borcea, C. (2017). Divert: A distributed vehicular traffic re-routing system for congestion avoidance. 16(1):58–72. DOI: 10.1109/TMC.2016.2538226.

Paranjothi, A., Khan, M. S., and Zeadally, S. (2020). A survey on congestion detection and control in connected vehicles. *Ad Hoc Networks*, 108:102277. DOI: 10.1016/j.adhoc.2020.102277.

Posner, J., Tseng, L., Aloqaily, M., and Jararweh, Y. (2021). Federated learning in vehicular networks: Opportunities and solutions. *IEEE Network*, 35(2):152–159. DOI: 10.1109/MNET.011.2000430.

Shen, X., Chen, J., Zhu, S., and Yan, R. (2024). A decentralized federated learning-based spatial–temporal model for freight traffic speed forecasting. *Expert Systems with Applications*, 238:122302. DOI: 10.1016/j.eswa.2023.122302.

Soto, I., Calderon, M., Amador, O., and Urueña, M. (2022). A survey on road safety and traffic efficiency vehicular applications based on c-v2x technologies. *Vehicular Communications*, 33:100428. DOI: 10.1016/j.vehcom.2021.100428.

Sun, P., Aljeri, N., and Boukerche, A. (2020). Machine learning-based models for real-time traffic flow prediction in vehicular networks. *IEEE Network*, 34(3):178–185. DOI: 10.1109/MNET.011.1900338.

Tennøy, A. and Hagen, O. H. (2021). Urban main road capacity reduction: Adaptations, effects and consequences. *Transportation Research Part D: Transport and Environment*, 96:102848. DOI: 10.1016/j.trd.2021.102848.

Treiber, M. and Kesting, A. (2013). *Traffic Flow Dynamics: Data, Models and Simulation*. Springer Berlin, Heidelberg. DOI: 10.1007/978-3-031-93922-8.

Xiao, H., Zhao, J., Pei, Q., Feng, J., Liu, L., and Shi, W. (2022). Vehicle selection and resource optimization for federated learning in vehicular edge computing. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):11073–11087. DOI: 10.1109/TITS.2021.3099597.

Ye, M., Fang, X., Du, B., Yuen, P. C., and Tao, D. (2023). Heterogeneous federated learning: State-of-the-art and research challenges. 56(3). DOI: 10.1145/3625558.

Yu, Z., Hu, J., Min, G., Xu, H., and Mills, J. (2020). Proactive content caching for internet-of-vehicles based on peer-to-peer federated learning. In *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, pages 601–608. DOI: 10.1109/ICPADS51040.2020.00083.

Yuan, L., Wang, Z., Sun, L., Yu, P. S., and Brinton, C. G. (2024). Decentralized federated learning: A survey and perspective. *IEEE Internet of Things Journal*, 11(21):34617–34638. DOI: 10.1109/JIOT.2024.3407584.

Zhang, R., Mao, J., Wang, H., Li, B., Cheng, X., and Yang, L. (2025). A survey on federated learning in intelligent transportation systems. *IEEE Transactions on Intelligent Vehicles*, 10(5):3043–3059. DOI: 10.1109/TIV.2024.3446319.

Zhang, S., Li, J., Shi, L., Ding, M., Nguyen, D. C., Tan, W., Weng, J., and Han, Z. (2024). Federated learning in intelligent transportation systems: Recent applications and open problems. *Trans. Intell. Transport. Sys.*, 25(5):3259–3285. DOI: 10.1109/TITS.2023.3324962.

Zhang, X., Liu, J., Hu, T., Chang, Z., Zhang, Y., and Min, G. (2023). Federated learning-assisted vehicular edge computing: Architecture and research directions. *IEEE Vehicular Technology Magazine*, 18(4):75–84. DOI: 10.1109/MVT.2023.3297793.

Zhao, J., Quan, H., Ge, P., Huang, Y., and Xiao, Y. (2025). Vehicular edge computing system: A survey. *IEEE Internet of Things Magazine*, pages 1–7. DOI: 10.1109/MIOT.2025.3590349.